DOCUMENT RESUME

ED 165 194 CS 502 249

AUTHOR Craig, Robert T.
TITLE Implications of Recent Research in the "Computational Paradigm."
PUB DATE Apr 78
NOTE 18p.; Paper presented at the Annual Meeting of the International Communication Association (Chicago, Illinois, April 25-29, 1978)

EDRS PRICE MF-$0.83 HC-$1.67 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Cognitive Processes; *Communication (Thought Transfer); Communication Skills; Comprehension; *Computational Linguistics; *Computer Science; Concept Formation; Content Analysis; Research; Theories
IDENTIFIERS *Communication Research

ABSTRACT

Describing cognitive science as a new discipline based upon the integration of elements of psychology, computer science, linguistics, philosophy, and education, this paper argues that such a discipline has important implications for the field of communication. Following an overview of the literature of cognitive science, the paper offers an explication of the computational paradigm underlying the new discipline. It then refutes attacks that have been made upon the paradigm's assumption that artificial intelligence research and psychological experimentation should be mutually heuristic because they both seek a unified theory of the cognitive processes performed by any system capable of certain tasks. In an illustration of the kinds of theoretical concepts that are emerging from cognitive science, the paper describes a recently proposed theory of story comprehension. In conclusion, the paper argues that cognitive science has important implications for communication research in the areas of content analysis methodology, the theory of messages, and the theory of communicative competence, and it outlines a proposed research project designed to illustrate those implications. (Author/FL)

ED165194

Implications of Recent Research in the "Computational Paradigm"

Robert T. Craig
Department of Speech Communication
The Pennsylvania State University
University Park, Pennsylvania 16802

A Paper presented to the annual convention of the International Communication Association, Information Systems Division program on Cognitive Representation; Chicago, Illinois; April 27, 1978.

The study of cognition has, in recent years, blossomed profusely in many fields. I report today on the work of a relatively small community of scholars who are attempting to integrate many disciplines concerned with cognition under the guidance of a new paradigm. Bobrow and Collins, in their 1975 edited volume, Representation and Understanding, announced the emergence of this new field, which they call "cognitive science."

> Cognitive science includes elements of psychology, computer science, linguistics, philosophy, and education, but it is more than the intersection of these disciplines. Their integration has produced a new set of tools for dealing with a broad range of questions. (Bobrow & Collins, 1975: ix-x)

The growth of cognitive science is reflected in the rapid appearance of new literature. The new journal Cognitive Science began publication in 1977. Relevant articles have appeared in journals such as The American Journal of Computational Linguistics, Artificial Intelligence, Cognition, Cognitive Psychology, and Memory and Cognition, among others. Major monographs and collections include Winograd (1972); Schank and Colby (1973); Bobrow and Collins (1975); Norman, Rummelhart and the LNR Research Group (1975); Schank (1975); Charniak and Wilks (1976); Schank and Abelson (1977); and Just and Carpenter (1977). Miller and Johnson-Laird (1976) and Anderson (1976), while not quite in the same group, are closely related works. This paper does not attempt to characterize this whole literature in any detail. Rather, I try to: 1) broadly describe the paradigm underlying cognitive science and the controversies surrounding that paradigm, 2) describe one example of a line of research which seems promising, and 3) suggest several implications of the new field for the field of communication.

Winograd (1977) suggests that the development of cognitive science represents a paradigm shift from a structural paradigm exemplified by Chomsky's work in linguistics to a computational paradigm focusing on

the study of cognitive processes. Although the computational paradigm is closely identified with fields, such as artificial intelligence and computational linguistics, in which electronic computers are heavily used, the new paradigm is not to be confused with the use of computers as such.

> The name 'computational' is not applied to this paradigm because computers are used in carrying out the research. One could imagine the concepts being developed without any direct use of computers, and a large percentage of the current applications of computers to the study of language do not fall within this paradigm at all. What is central is the metaphor provided by viewing human cognitive capacity as a kind of 'physical symbol system', and drawing parallels between it and those physical symbol systems we are learning to construct out of electronic components. The parallels are not at the level of the physical components, but at the level of the abstract organization of processes and symbol structures. (Winograd, 1977: 168-169)

The computational paradigm, then, generates research questions such as these: How is discourse understood for purposes such as recalling, paraphrasing, summarizing, answering questions, translating into foreign languages, etc.? What knowledge structures are required for such complex symbolic processes? Questions such as these are investigated by a mix of methodologies including conceptual analysis, computer programming and psychological experimentation. The enterprise emphasizes the writing of computer programs which perform the tasks under investigation. For example, consider attempting to write a computer program that would read a newspaper story and produce a paraphrase of that story. This is more than just a difficult programming problem; it is a difficult programming problem because it is first of all a difficult theoretical problem, or rather several of them: What cognitive structures are needed to represent the information contained in the story? By what means can the linguistic text be translated into those structures? What linguistic and extra-linguistic

knowledge is needed to extract the key facts from the story? What processes can be used to bring that knowledge to bear effectively? In short, the computational problem of how to paraphrase a newspaper story becomes a rich source of relatively well-defined theoretical problems. The computational paradigm assumes that the solution of those problems will help to explain the operation of "physical symbol systems" generally. To the extent that the resulting theoretical ideas are implemented in working computer programs, cognitive science claims to have constructed a working model of the relevant cognitive processes, and thus to have achieved some degree of understanding of them. Psychological experimentation can attempt to discover the extent to which humans perform the same processes. Psychological data, too, are regarded as relevant to theory construction and should constrain the design of computer algorithms. Computer programming and psychological experimentation should be mutually heuristic because they seek the same end: a unified theory of the cognitive processes performed by any system which is capable of performing certain tasks.

These assumptions have been severely criticized. Weizenbaum (1976) has argued that artificial intelligence research is morally reprehensible because it fosters an essentially mechanistic model of human nature, especially of the higher mental faculties. He also argues that research following the computational paradigm has been mostly technological rather than theoretical, and thus is not genuinely scientific. Dresher and Hornstein (1976) have made a better-documented case for the same charge of theoretical vacuity. Anderson (1976), in his basically sympathetic comments, argues convincingly that the strongest claims of the abstract similarity of mental and computational processes are essentially unprovable.

In my view, the outcome of the whole debate hinges on the question of scientific versus technological--theoretical versus atheoretical--status of the computational paradigm. Weizenbaum's moral argument plucks a responsive chord but finally does not convince. The same kinds of moral judgments have been made against many scientific advances from Galileo to the present. Aerodynamics is a theoretical problem; how to build an airplane is a technological problem; whether "man was meant to fly" is a rather uninteresting moral problem. The moral meaning of computer programs that mimic the mental skills of people, no less than the moral meaning of the airplane, is a matter of cultural definition and need not threaten the ethical status of human nature once assimilated into a culture which rationally affirms that status. Knowledge may banish innocence but by itself cannot engender evil, or so I believe.

The matter of scientific status is more difficult to judge in the abstract. Weizenbaum, himself a prominent contributer to artificial intelligence research in earlier days, charges his colleagues with the sin of hubris. His case largely rests on extreme quotations and outdated examples. Dresher and Hornstein make a stronger, but ultimately unconvincing case. As Winograd (1977) points out in his reply to their paper, Dresher and Hornstein argue from the somewhat narrow paradigmatic confines of the Chomskian school of transformational-generative (TG) linguistics, which has goals very different from those of the computational paradigm. The critique draws some blood with two points: that artificial intelligence researchers often have not clearly distinguished theoretical from merely technical computational issues, and that Schank and Winograd, among others, have done some sloppy linguistic reasoning in constructing certain parts of their systems. These arguments can

be turned around, however. Is not the grammarian's fetish for formal elegance as distracting from substantive theoretical concerns as is the artificial intelligence weakness for computer programming tricks? I cannot see how mere formal economy speaks directly to what Dresher and Hornstein regard as "the central problem of a theory of language," which is "to explain how people are able to learn their native language." (323) Dresher and Hornstein are at pains to exempt TG theory from the often-made charge that it denies the importance of cognitive processing. TG theory, they say, is not intended to be a complete theory of language use, which "must specify, in addition to a theory of competence, a theory of production and a theory of parsing" (378). They assert that a TG theory of linguistic competence is a necessary prerequisite for a theory of parsing, however, because it has to function as a template against which to match that which is parsed. This argument, which beams with presumptive reasonableness, displays most clearly the paradigmatic myopia of Dresher and Hornstein. If any theoretical generalization has emerged from cognitive science, it is that linguistic competence in the TG sense is a prerequisite for only certain special cases of discourse comprehension. For the most part, semantic and nonlinguistic world knowledge, supplemented by some minimal syntactical resources, are sufficient. How else is it that we understand normal conversational language, which from the grammatical point of view is hopelessly degenerate? What linguistic competence is actually required for parsing is a matter to be discovered in the empirical and theoretical study of parsing, not to be arbitrated by the formal passions of the grammarians. Dresher and Hornstein are correct in asserting that the computational paradigm has not yet produced a coherent theory of cognitive processes. TG grammar, one might retort, has not produced

a fully coherent theory of competence. More to the point, I see in recent books such as Bobrow and Collins (1975) and Schank and Abelson (1977), the beginnings of real theoretical substance in cognitive science. I will illustrate this comment in a moment.

I turn briefly to the critical comments of Anderson (1976). Anderson presents a logical proof that it is not possible to determine a unique theory to account for the input-output relations which are the data of cognitive psychology. Indeed, there are an indefinitely large number of equivalent theories. This point tells against those who are "committed to arguing that at an abstract level, the human and the computer are the same sort of devices. This may be true but the problem is that a unique abstract characterization of man's cognitive functioning does not exist. So it is somewhat pointless to try to decide what kind of device he is" (Anderson, 1976: 15). Yet Anderson maintains that the experience of cognitive psychology demonstrates that "the best way to theorize is by proposing a model of internal structure and process" (15). Actually, given the mind as a black box which performs complex tasks, we have little choice but to follow this strategy. In fact, Anderson's argument could be taken as a vindication rather than an indictment of the computational approach, for it suggests that abstract models which posit representations and processes *sufficient* to account for the competencies underlying complex input-output relations are the best that any science of cognition could ever hope to achieve. An additional point, which will come up again later in this paper, is that what is an unfortunate limitation of cognitive science from the standpoint of psychology may not matter at all from the standpoint of a discipline of communicology focused on the study of messages, a discipline which would care about the representation of

message content but would care not what is in the black box of the mind.

To summarize this first section of the paper, the new field of cognitive science, which has grown explosively in recent years, rests on a coherent assumptive base, the computational paradigm. The paradigm has been attacked on various grounds, but, in my view, those attacks have been unsuccessful. The ultimate test of the paradigm will be the degree of scientific progress achievable within it.

To illustrate the kinds of theoretical concepts that are emerging in cognitive science, I now describe an example that I have found particularly interesting: the theory of story understanding presented by Schank and Abelson (1977).

Schank and Abelson's basic purpose is to characterize the knowledge that is required to understand connected discourse and how that knowledge is used in various tasks which require understanding. Focusing on story narratives, they try to describe the structures and processes required to read a story and perform such tasks as 1) expand the story into a longer, more detailed story, 2) abstract key ideas and summarize the story, 3) answer questions about the story (not unlike the questions on a reading comprehension test), or 4) translate the story into a foreign language. Associates of the authors have written computer programs which perform those tasks at a fairly sophisticated level within limited but real domains. The book is an account of the theoretical insights which those programs reflect and which have been stimulated by the problems encountered in writing such programs.

The most basic theoretical move in this work is to unite content and form--knowledge and inference procedures--in a single system of representations. This is a departure from mathematically oriented artificial intelligence research, which has tended to separate form

and content into a knowledge base and a set of powerful logical mechanisms which operate on that base. Such systems, despite their formal elegance, run up against insuperable problems of complexity when confronted by realistic discourse-understanding tasks. Schank and Abelson attempt to be "more pragmatic about knowledge" (3). They present a direct challenge to cognitive theory:

> There is a very long theoretical stride...from the idea that highly structured knowledge dominates the understanding process, to the specification of the details of the most appropriate structures. It does not take one very far to say that schemas are important: one must know the content of the schemas. To be eclectic here is to say nothing. If one falls back on the abstract position that only form is important, that the human mind is capable of developing knowledge structures of infinitely varied content, then one sacrifices the essence of the structure concept, namely the strong expectations which make reality understandable. In other words, a knowledge structure theory must make a commitment to particular content schemas. (Schank and Abelson, 1977: 10)

The knowledge structures which Schank and Abelson have devised for understanding narratives are presented as six roughly heirarchical levels. I say "roughly heirarchical" because the levels call upon each other in complex ways during processing. The levels, listed from less to more general, are: conceptualizations, causal chains, scripts, plans, goals and themes.

<u>Conceptualizations</u> function at the linguistic level to map sentences into a scheme of general concepts. Structures at this level allow the understanding system to recognize paraphrases--different sentences which are similar in meaning. Schank's Conceptual Dependency Theory asserts that all acts described in story narratives can be expressed in terms of eleven primitive acts. To briefly illustrate the functioning of a system at this level, consider the two sentences:

(1) John gave $1 to the cashier.

(2) The cashier took $1 from John.

These sentences would be mapped into conceptualizations such as:

(1) John ATRANS $1 from John to cashier.

(2) Cashier ATRANS $1 from John to cashier.

This representation highlights the primitive concept underlying both give and take in this context (ATRANS is one of Schank's eleven primitives), and also highlights the fact that the sentences overlap in meaning except for the identity of the acting agent.

Causal chains, the second level of representation, detect linguistic ties between events. For example, consider this little story:

(3) The alarm rang at 7:00 a.m.. The ringing woke up Bob. Then Bob got up and took a shower.

The causal chaining of the three sentences is created by obvious linguistic devices.

Narratives, however, often do not linguistically cue the connections among events; yet we have no trouble understanding them. How is this accomplished? Schank and Abelson have devised a third level of representation, called scripts, to account for this. Consider the following two stories:

(4) John went to a restaurant. He asked the waitress for a coq au vin. He paid the check and left. (p. 38)

(5) John was walking on the street. He thought of cabbages. He picked up a shoe horn. (p. 39)

The difference in comprehensibility between these two stories results from the fact that story (4) refers to a standard sequence of events associated with a certain setting, in this case a restaurant, while story (5) does not. Such a conventional sequence of events is called a script. Knowing an applicable script enables the story understander to fill in missing details and to recognize the whole structure as orderly.

"Not all sequences of events are conventional, however. The following story must be understood by some process other than causal chains and scripts:

(8) Willa was hungry. She took out the Michelin Guide. (p. 71)

Schank and Abelson claim that story understanders have more general knowledge structures, called plans, which state standard means for achieving certain general goals. Here the second sentence must be understood as a means employed to satisfy the need of hunger. The hypothetical understanding process is complex, but it generally involves recognizing that one source of food is a restaurant, that to get to a restaurant one has to know where it is, that one source of information is a book, and that the Michelin Guide is a book. Thus the story is understandable in terms of plans.

The fifth level of representation, called goals, was devised to deal with stories which involve plans but in which the goal of an actor is not explicit. Consider this story (condensed from Schank and Abelson's version, p. 106):

(9) John...rushed to the theatre, only to find out that all the tickets were sold. He decided to go to a concert in Brooklyn.

This story is comprehensible because we recognize that plays and concerts are both forms of entertainment and so are reasonably substitutable for each other. Schank and Abelson have devised an extensive classification of what they call Goal Forms. This knowledge base is used in understanding stories by recognizing the goals of actors.

Sometimes, however, goals cannot be recognized without recourse to more general knowledge about the actors involved. Consider the difference between these two events:

(10) George the garbage man picked up the garbage from all the cans on the street.

(11) Larry the lawyer picked up the garbage from all the cans on the street. (p. 133)

Themes, the sixth and final level of representation, consist of background information which allows a story understander to predict that a certain actor would have a certain goal. Garbage men are more likely to collect garbage than are lawyers. There are three kinds of themes: 1) Role Themes like lawyer and garbage man, 2) Interpersonal Themes like lover and friend, and 3) Life Themes like success and do-good-works. Recognition of themes is essential to understanding many narrative structures.

The preceding pages have been a brief, but I hope, helpful attempt to illustrate one kind of theory beginning to emerge from cognitive science. I now turn to discuss the implications of cognitive science for communication theory.

Recent work in the computational paradigm may have important implications for many areas of communication research. Here I would like to draw implications in three areas: content analysis methodology, the theory of messages, and the theory of communicative competence.

First, computational models of cognitive representation could have a major impact on the method of content analysis. Krippendorff (1969) introduced the concept of "models of messages" to refer to the analytical constructs which underlie content analysis procedures. Models of messages are, in effect, models of cognitive representation in a communication system. Current content analysis computer programs, such as The General Inquirer, are essentially word-counting programs which imply what Krippendorff would call an "association model" of

messages, a model which sees the message as a sequence of behaviors (words) produced according to probabilistic laws of association. The next generation of content analysis computer programs, drawing upon current work in artificial intelligence, could be based upon what Krippendorff calls "discourse models of messages", which treat the message as having been generated from cognitive structures in the source. The computer programs would extract ideas from messages rather than merely count certain classes of words.

The second implication is for a theory of messages. Much of the work in the computational paradigm, but especially the recent work of Schank and Abelson, has an interesting quality which has led me to reflect on the perennial question of whether there could be a discipline of communicology with a unique central focus. One often hears that the central focus would be the study of messages; yet the field of communication seems always to be overpowered by centrifugal forces which throw its various parts into psychology, sociology, political science, linguistics, and yes, even computer science. This is not a new problem. Aristotle's book on rhetoric could be accused of being a mash of logic, psychology, ethics, politics and poetics. Since he wrote other books on all those subjects, one might wonder why a book on rhetoric was necessary. In fact, however, Aristotle's rhetoric is a book about a unique subject: how to decide what to say in public speaking situations. The logic, psychology and politics contained in rhetoric are very special versions of those subjects--really, not those subjects at all, but what corresponds to those subjects in the forms of popular discourse as opposed to philosophy. Oddly enough, Schank and Abelson's work on the computer understanding of narratives strongly reminds me of this aspect of Aristotle's rhetoric. Their

13

theory of conceptualizations, scripts, plans, goals and themes comprises a description of part of the world of human action, but not a truly psychological description. The theory is not a theory of human action but, rather, a theory of what one has to assume about human action in order to comprehend ordinary stories involving action. Now this is a beginning move toward what I would call a theory of messages. It strikes me that this kind of theory could become the central focus of at least an important subdiscipline of communicology concerned with understanding message processes. This would not be a theory about what is in the black box of the mind but a theory of what people say to each other.

Finally, I believe that research in the computational paradigm could be of significant assistance in building a theory of communicative competence. What I have in mind here is a specific research project which sounds so good that I plan to invest part of the next summer in exploring its feasibility. The project is to build a model of interpersonal conversation. The key idea is to start with the most primitive model possible and to systematically add to that basic model those competencies required to produce realistic conversational behavior. The basic model would comprise a set of subroutines called conversationalists, each composed of 1) a set of speeches or things to say, and 2) a set of triggering conditions which determine which speech is selected next. This model accords with Goffman's notion that most conversation involves running off tapes of past experiences (Goffman, 1974: 504). The model also permits us to start at a near-zero level of communicative competence at which each conversationalist simply rattles off its speeches according to a predetermined order of priority, totally unresponsive to the speeches of the other conversationalist. What competencies have to be built into this trivial model in order to

make it perform realistically? To begin with, the basic model may not always be terribly unrealistic, for conversationalists are often not as responsive to each other as one might suppose; conversation is a loosely structured affair (Goffman, 1974: 502). Responsiveness, which would be reflected in the appropriateness of speech-selection, is a matter of degree and a cumulative result of several skills (Goffman, p. 500, gives one list of conversational skills). One competency might involve a topicality device such that the conversationalist would recognize key terms in the speeches of the partner and select speeches from its own repertory relating to the same topics. This perhaps fairly simple device might add a great deal of realism to the simulation by generating a more orderly development of ideas. Another competency might involve recognition of an episode structure, so that opening and closing moves would be produced as required, or recognized and responded to appropriately. Another competency would involve more sophisticated turn-taking behavior based on each conversationalist's "continuous" monitoring of the other's verbal and "nonverbal" output. The system, at this point, lacks all but the most primitive linguistic knowledge. Additional linguistic competence would permit still more coherent idea development and would permit the conversationalist's speeches to be adapted to the conversational, social and "physical" contexts. The question is, what minimal linguistic competencies are necessary to produce competent conversational behavior? I suspect that the most crucial competencies will prove to be pragmatic ones having to do with a logic of speech acts in the conversational context (Frentz and Farrell, 1976). Still other competencies would involve self-disclosure strategies, patterns of affect and dominance and other socioemotional dimensions of behavior. My ideas as to how to construct

these competencies are still very ill-formed. The point, however, is just that the problem of building an artificial conversationalist structures the theoretical problem of communicative competence in a potentially very fruitful way. This is what I see as the exciting promise of the computational paradigm.

References

ANDERSON, JOHN R. Language, Memory and Thought. Hillsdale, N.J.: Laurence Erlbaum Associates, 1977.

BOBROW, DANIEL G. & COLLINS, ALLAN (Eds). Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.

CHARNIAK, EUGENE & WILKS, YORICK (Eds). Computational Semantics: An Introduction to Artificial Intelligence and Natural Language Comprehension. New York: North-Holland Publishing Company, 1976.

DRESHER, B. ELAN & HORNSTEIN, NORBERT. On some supposed contributions of artificial intelligence to the scientific study of language. Cognition, 4 (1976) 321-398.

GOFFMAN, ERVING. Frame Analysis. New York: Harper & Row, 1974.

JUST, MARCEL ADAM & CARPENTER, PATRICIA A. (Eds). Cognitive Processes in Comprehension. Hillsdale, N.J.: Laurence Erlbaum Associates, 1977.

KRIPPENDORFF, KLAUS. Models of Messages: Three Prototypes. Chapter 5 in George Gerbner et al. (Eds) The Analysis of Communication Content. New York: John Wiley & Sons, 1969.

MILLER, GEORGE A. & JOHNSON_LAIRD, PHILLIP N. Language and Perception. Cambridge, Mass.: Belknap Press (Harvard Univ. Press), 1976.

NORMAN, DONALD A. & RUMELHART, DAVID E. & the LNR RESEARCH GROUP. Explorations in Cognition. San Francisco: W.H. Freeman, 1975.

SCHANK, ROGER C. Conceptual Information Processing. New York: American Elsevier Publishing Company, 1975.

SCHANK, ROGER C. & ABELSON, ROBERT B. Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures. Hillsdale, N.J.: Laurence Erlbaum Associates, 1977.

SCHANK, ROGER C. & COLBY, KENNETH MARK (Eds). Computer Models of Thought and Language. San Francisco: W.H. Freeman, 1973.

WEIZENBAUM, JOSEPH. Computer Power and Human Reason: From Judgment to Calculation. San Francisco: W.H. Freeman, 1976.

WINOGRAD, T. Understanding Natural Language. N.Y.: Academic Press, 1972.

WINOGRAD, T. On some contested suppositions of generative linguistics about the scientific study of language. Cognition, 5(1977)151-179.

Abstract of: Robert T. Craig, "Implications of Recent Research in the 'Computational Paradigm'" a paper presented to the annual convention of the International Communication Association, Information Systems Division program on Cognitive Representation, Chicago, Illinois, April 27, 1978.

This paper describes the new, interdisciplinary field of cognitive science and argues that the new field has important implications for the communication discipline. The rapidly growing literature of cognitive science is overviewed and the "computational paradigm" underlying the new field is explicated. The computational paradigm assumes that artificial intelligence research and psychological experimentation should be mutually heuristic because they seek the same end: a unified theory of the cognitive processes performed by any system which is capable of performing certain tasks. The assumptions of the computational paradigm have been attacked by several critics. The paper argues that those attacks have been unsuccessful. To illustrate the kinds of theoretical concepts which are emerging from cognitive science, the paper describes in some detail the theory of story comprehension recently proposed by Roger Schank and Robert Abelson. Finally, the paper argues that cognitive science has important implications for communication research in three areas: content analysis methodology, the theory of messages and the theory of communication competence. A proposed research project illustrates those implications.